arXiv:1511.04001v2 [q-bio.GN] 20 Jul 2016

# *rasbhari*: Optimizing Spaced Seeds for Database Searching, Read Mapping and Alignment-Free Sequence Comparison

Lars Hahn[1], Chris-André Leimeister[1], Rachid Ounit[2], Stefano Lonardi[2], and Burkhard Morgenstern[1,3]

[1] University of Göttingen, Department of Bioinformatics, Goldschmidtstr. 1, 37077 Göttingen, Germany
[2]University of California, Riverside, Department of Computer Science and Engineering, 900 University Ave. Riverside, CA 92521, USA
[3]University of Göttingen, Center for Computational Sciences, Goldschmidtstr. 7, 37077 Göttingen, Germany

July 21, 2016

**Abstract**

Many algorithms for sequence analysis rely on word matching or word statistics. Often, these approaches can be improved if binary patterns representing *match* and *don't-care* positions are used as a filter, such that only those positions of words are considered that correspond to the *match positions* of the patterns. The performance of these approaches, however, depends on the underlying patterns. Herein, we show that the *overlap complexity* of a pattern set that was introduced by Ilie and Ilie is closely related to the *variance* of the number of matches between two evolutionarily related sequences with respect to this pattern set. We propose a modified hill-climbing algorithm to optimize pattern sets for database searching, read mapping and alignment-free sequence comparison of nucleic-acid sequences; our implementation of this algorithm is called *rasbhari*. Depending on the application at hand, *rasbhari* can either minimize the *overlap complexity* of pattern sets, maximize their *sensitivity* in database searching or minimize the *variance* of the number of pattern-based matches in alignment-free sequence comparison. We show that, for database searching, *rasbhari* generates pattern sets with slightly higher *sensitivity* than existing approaches. In our *Spaced Words* approach to

alignment-free sequence comparison, pattern sets calculated with *rasbhari* led to more accurate estimates of phylogenetic distances than the randomly generated pattern sets that we previously used. Finally, we used *rasbhari* to generate patterns for short read classification with *CLARK-S*. Here too, the sensitivity of the results could be improved, compared to the default patterns of the program. We integrated *rasbhari* into *Spaced Words*; the source code of *rasbhari* is freely available at `http://rasbhari.gobics.de/`

## Author Summary

We propose a fast algorithm to generate spaced seeds for database searching, read mapping and alignment-free sequence comparison. Spaced seeds – *i.e.* patterns of match and don't-care positions – are used by many algorithms for sequence analysis; designing optimal seeds is therefore an active field of research. In sequence-database searching, one wants to optimize sensitivity, *i.e.* the probability of finding a region of homology; this can be done by minimizing the so-called overlap complexity of pattern sets. In alignment-free DNA sequence comparison, the number $N$ of pattern-based matches is used to estimate phylogenetic distances. Here, one wants to minimize the variance of $N$ in order to obtain stable phylogenies. We show that for spaced seeds, the overlap complexity – and therefore the sensitivity in database searching – is closely related to the variance of $N$. Our algorithm can optimize the sensitivity, overlap complexity or the variance of $N$, depending on the application at hand.

## Introduction

$k$-mers, *i.e.* words of length $k$, are used in many basic algorithms for biological sequence comparison. Word matches are used, for example, as *seeds* in the *hit-and-extend* approach to database searching and read mapping [2, 41, 17]. In alignment-free sequence comparison, sequences are represented as word-frequency vectors to estimate distances or similarities between them, *e.g.* as a basis for phylogeny reconstruction [9, 43, 49, 22, 1], see [48, 3] for reviews. Similarly, word statistics are used to classify DNA or protein sequences [26, 39, 32], for datamining [33] and for remote homology detection [29]. It is well known that many word-based approaches produce better results if *spaced words* or *seeds* are used instead of the previously used *contiguous* words or word matches. That is, for a pre-defined binary pattern $P$ representing *match* and *don't-care* positions, one considers only

those positions in a word of the same length that correspond to the *match* positions of $P$.

Pattern-based word matching has been proposed for *hit-and-extend* database searching by Ma *et al.* [30], see also [8]; *multiple* spaced seeds are now a standard *filtering* step in homology searching [15, 6]. Spaced seeds are also routinely used in metagenome sequence clustering and classification [7, 38], protein classification [37], read mapping [40, 35], to find *anchor points* for multiple sequence alignment [11, 10] and for alignment-free phylogeny reconstruction [24]; similarly, the *average common substring* approach to sequence comparison [47] could be improved by allowing for mismatches [16, 25, 45, 46, 44]. Brejova *et al.* extended the concept of *spaced seeds* to homologies among protein-coding regions [4] and introduced *vector seeds* [5]. In general, the advantage of pattern-based approaches is the fact that spaced-word occurrences at neighbouring sequence positions are statistically less dependent than occurrences of contiguous words [28, 34].

In pattern-based approaches, the underlying patterns of match and don't-care positions are of crucial importance for the quality of the results. Generally, non-periodic patterns are preferred since they minimize redundancies between overlapping matches and lead to a more even distribution of matches. This increases the probability of obtaining a hit between two homologous sequences in database searching and leads to more stable distance estimates in phylogeny reconstruction. Noé and Martin [36] defined a *coverage criterion* for multiple spaced seeds and showed that this criterion is related to the *Hamming distance* between two sequences. In the context of database searching, patterns or sets of patterns are often called *seeds*. (Originally, the word *seed* denoted a match of – contiguous or spaced – words between a query and a database sequence that could be extended to the left and to the right. But now *seed* often denotes the underlying pattern in pattern-based approaches).

In database searching, one wants to maximize the *sensitivity* of pattern sets *i.e.* the probability of finding at least one hit within a gap-free alignment of a given length $L$ and probability $p$ for a match between two residues. Calculating the sensitivity of a pattern set is *NP-hard* [28]. The sensitivity can be approximated by dynamic programming [30, 27], but the run time of this algorithm is still exponential in the length of the pattern. In *PatternHunter II*, a *greedy* algorithm is used to find suitable patterns. In 2007, Ilie and Ilie introduced the *overlap complexity* of a pattern set and showed experimentally that – for a given number of patterns with a given length and number of match positions – minimizing the overlap complexity corresponds to maximizing the *sensitivity* in database searching [20]. In contrast

to the sensitivity, however, the overlap complexity can be easily calculated. To find optimal pattern sets, Ilie and Ilie proposed a *hill-climbing algorithm* that minimizes the overlap complexity. They implemented their algorithm in a software tool called *SpEED* [21], which is several orders of magnitude faster than competing approaches and is now considered the state-of-the-art in seed optimization.

Recently, we proposed to use *spaced-word* frequencies instead of word frequencies for alignment-free sequence comparison [24, 19]. We showed that phylogenetic trees calculated from spaced-word frequencies are more accurate than trees calculated from contiguous-word frequencies. As in database searching, our results could be improved by using *multiple* patterns. In our original study, we used randomly generated multiple patterns of *match* and *don't-care* positions. In a follow-up paper, we studied the number $N$ of spaced-word matches between two DNA sequences for a set of binary patterns [34]. Our data suggest that minimizing the variance of $N$ for pattern sets improves alignment-free phylogeny reconstruction.

In this paper, we first show that the variance of the number $N$ of spaced-word matches is closely related to the *overlap complexity* of the underlying set of patterns. We propose a modified hill-climbing algorithm that can be used to generate pattern sets, either with minimal variance of $N$, or with minimal overlap complexity, or with maximal sensitivity in database searching, depending on the application at hand. While the algorithm proposed in [20] iterates over all patterns $P$ in a set $\mathcal{P}$ of patterns and all pairs of positions in $P$ to improve $\mathcal{P}$, we calculate for each pattern $P \in \mathcal{P}$ how much $P$ contributes to the variance or overlap complexity, respectively, of $\mathcal{P}$. We then modify those patterns first that contribute most to the variance or complexity.

The implementation of our approach is called *rasbhari (Rapid Approach for Seed optimization Based on a Hill-climbing Algorithm that is Repeated Iteratively)*. Experimental results show that pattern sets calculated with *rasbhari* have a slightly higher sensitivity in database searching than pattern sets calculated with *SpEED*, while the run time of both programs is comparable. In alignment-free sequence comparison, we obtain more accurate phylogenetic distances if we use *rasbhari* to minimize the variance of $N$ for the underlying pattern sets, than we obtained with the randomly generated pattern sets that we previously used. In a third application, we used pattern sets generated with *rasbhari* in the program *CLARK-S* [38] for short read classification. The sensitivity of the classification could be improved in this way, while *rasbhari* is substantially faster than the method that is used by default for pattern generation in *CLARK-S*.

# Methods

## Overlap complexity

We consider sets $\mathcal{P} = \{P_1, \ldots, P_m\}$ of binary patterns, where $\ell_r$ is the length of pattern $P_r$ and $\ell = \max_r \ell_r$. That is, each $P_r$ is a word of length $\ell_r$ over the alphabet $\{1, 0\}$. A '1' in a pattern $P_r$ represents a *match* position, a '0' a *don't-care* position. For a single pattern $P_r$, the number of match positions is called its *weight* $w$. For simplicity, we assume that all patterns in a set $\mathcal{P}$ have the same *weight*.

In [34], we considered for two patterns $P_r, P_{r'}$ and $s \in \mathbb{Z}$ the number $n(P_r, P_{r'}, s)$ of positions that are match positions of $P_r$ *or* match positions of $P_{r'}$ (or both), if $P_{r'}$ is shifted by $s$ positions to the right, relative to $P_r$. If $s$ is negative, $P_{r'}$ is shifted to the left. For $P_r = 101011, P_{r'} = 111001$, for example, if $P_{r'}$ is shifted by 2 positions to the right, relative to $P_r$, then there are 6 positions (marked by asterisks below) that are match positions of $P_r$ or $P_{r'}$. Thus, for $s = 2$, we have $n(P, P_{r'}, 2) = 6$:

$$\begin{array}{lcccccccc} P_r: & 1 & 0 & 1 & 0 & 1 & 1 & & \\ P_{r'}: & & & 1 & 1 & 1 & 0 & 0 & 1 \\ & * & & * & * & * & * & & * \\ & & & \$ & & \$ & & & \end{array}$$

For the same situation, Ilie and Ilie [20] defined $\sigma[s] = \sigma_{r,r'}[s]$ as the number of positions where $P_r$ *and* $P_{r'}$ have a match positions, such as the positions marked by '$' above. In the above example one would therefore have $\sigma[2] = 2$. The *overlap complexity (OC)* of a set of patterns $\mathcal{P} = \{P_1, \ldots, P_m\}$ is then defined in [20] as

$$\sum_{r \le r'} \sum_{s=1-\ell_{r'}}^{\ell_r - 1} 2^{\sigma_{r,r'}[s]} \tag{1}$$

Note that, since for any two patterns $P_r, P_{r'}$ and $s \in \mathbb{Z}$, the equality

$$\sigma_{r,r'}[s] = 2w - n(P_r, P_{r'}, s)$$

holds, the *overlap complexity* of a set $\mathcal{P}$ can be written as

$$\sum_{r \le r'} \sum_{s=1-\ell_{r'}}^{\ell_r - 1} 2^{\sigma_{r,r'}[s]} \;=\; 2^{2w} \cdot \sum_{r \le r'} \sum_{s=1-\ell_{r'}}^{\ell_r - 1} (1/2)^{n(P_r, P_{r'}, s)} \tag{2}$$

Consequently, if we are looking at sets $\mathcal{P}$ of $m$ patterns with fixed weight $w$ and lengths $\ell_r$, then minimizing the overlap complexity of $\mathcal{P}$ is equivalent to minimizing the sum

$$\sum_{r \le r'} \sum_{s=1-\ell_{r'}}^{\ell_r - 1} (1/2)^{n(P_r, P_{r'}, s)} \tag{3}$$

Ilie and Ilie showed experimentally that the *OC* is closely related to the *sensitivity* of a pattern set. More precisely, they showed that for pattern sets with a given number of patterns of given lengths and weight, minimizing the *OC* practically amounts to maximizing the *sensitivity*. Consequently, in order to find suitable pattern sets for hit-and-extend approaches in database searching, they proposed to search for pattern sets with minimal *OC*. The main advantage of this approach is the fact that the *OC* of a pattern set is much easier to calculate than its *sensitivity*.

### Variance of the number of spaced-word matches

For a pattern $P$ of length $\ell$, we say that two sequences $S_1$ and $S_2$ have a *spaced-word match* with respect to $P$ at $(i, j)$, if the $\ell$-mers starting at $i$ and $j$ have identical characters at all *match* positions of $P$, *i.e.* if one has $S_1(i + \pi - 1) = S_2(j + \pi - 1)$ for all match positions $\pi$ in $P$. The sequences below, for example, have a spaced-word match at $(2, 4)$ with respect to the *pattern* $P = 110101$. Indeed, the 6-mers starting at positions 2 and 4 of the sequences are identical at all positions corresponding to a *match position* ('1') in $P$, while positions at *don't-care positions* ('0') may be matches or mismatches.

| | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| $S_1$ : | | | $A$ | $A$ | $T$ | $C$ | $G$ | $A$ | $T$ | $C$ | $A$ |
| $S_2$ : | $C$ | $G$ | $T$ | $A$ | $T$ | $T$ | $G$ | $A$ | $T$ | $T$ | |
| P: | | | | 1 | 1 | 0 | 1 | 0 | 1 | | |

In [34], we considered spaced-word matches between two sequences $S_1$ and $S_2$ with respect to a set $\mathcal{P} = \{P_1, \ldots, P_m\}$ of patterns, so-called $\mathcal{P}$-matches. Note that there can be up to $m$ $\mathcal{P}$-matches at each pair $(i, j)$ of positions of $S_1$ and $S_2$, one $\mathcal{P}$-match for each pattern $P_r$ in $\mathcal{P}$. We studied the number $N = N(S_1, S_2, \mathcal{P})$ of $\mathcal{P}$-matches between sequences $S_1$ and $S_2$ under a simplified model of evolution without insertions and deletions, with a *match probability* $p$ for pairs of homologous positions and a *background* match probability of $q$. It is easy to see that, for a pattern set $\mathcal{P}$, the *expected* number of $\mathcal{P}$-matches depends only on the number $m$ of patterns in $\mathcal{P}$ and

on their lengths $\ell_i$ and their weight $w$, *i.e.* number of match positions, but not on the particular sequence of *match* and *don't-care* positions in $\mathcal{P}$. The *variance* of $N$, however, does depend on the sequence of *match* and *don't-care* positions.

As discussed in [34], many alignment-free distance or similarity measures are – explicitly or implicitly – a function of the number $N$ of (spaced) word matches. To obtain stable distance measures for phylogeny reconstruction, it is therefore desirable to use pattern sets with a low variance of $N$. For a given set $\mathcal{P} = \{P_1, \ldots, P_m\}$ of patterns of lengths $\ell_1, \ldots, \ell_m$ and weight $w$, and with the above simple model of evolution, the variance of $N$ can be approximated by

$$\begin{aligned} Var(N) \approx (L-\ell+1) \cdot \sum_{r \le r'} \sum_{s \in R(r,r')} \left( p^{n(P_r, P_{r'}, s)} - p^{2w} \right) \\ + \quad (L-\ell+1) \cdot (L-\ell) \cdot \sum_{r \le r'} \sum_{s \in R(r,r')} \left( q^{n(P_r, P_{r'}, s)} - q^{2w} \right) \end{aligned} \tag{4}$$

where $L$ is the length of $S_1$ and $S_2$, respectively, and

$$R(r,r') = \begin{cases} \{1-\ell_{r'}, \ldots, \ell_r - 1\} & \text{if } r < r' \\ \{0, \ldots, \ell_r - 1\} & \text{if } r = r' \end{cases}$$

is the *range* in which $P_{r'}$ is to be shifted against $P_r$ [34]. Note that for different patterns $P_{r'} \neq P_r$ we have to consider all shifts between $1-\ell_{r'}$ and $\ell_r - 1$ of $P_{r'}$ against $P_r$ , for example:

$$\begin{array}{rccccccccccccccccccc} P_r: & & & & & 1 & 0 & 1 & 1 & & & & & 1 & 0 & 1 & 1 & & & \\ P_{r'}: & 1 & 0 & 1 & 0 & 1 & & & & , & \cdots & , & & & & & 1 & 0 & 1 & 0 & 1 \\ s: & & & -4 & & & & & & & & & & & & & & & 3 & & \end{array}$$

By contrast, if a pattern $P_r$ is shifted against itself, only shifts between 0 and $\ell_r - 1$ need to be considered, to avoid double counting of shifts[1], for example:

$$\begin{array}{rccccccccccccccc} P_r: & 1 & 0 & 1 & 1 & & & & & 1 & 0 & 1 & 1 & & & \\ P_r: & 1 & 0 & 1 & 1 & , & \cdots & , & & & & & 1 & 0 & 1 & 1 \\ s: & & & 0 & & & & & & & & & & 3 & & \end{array}$$

On the right-hand side of (4), the first summand is the variance of the 'homologous' spaced-word matches (in a model without indels, these are

[1] In [34], we ignored this fact and gave a slightly different estimate for $Var(N)$.

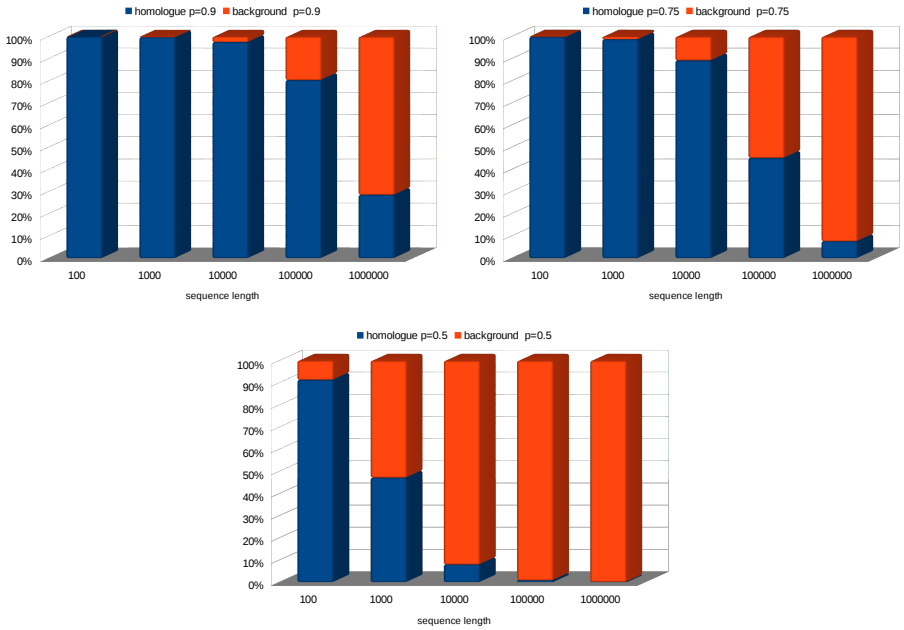


Figure 1: Homologue and background variance Contribution of the *homologue* and *background* variance to the total variance of the number $N$ of spaced-word matches in equation (4) for different match probabilities $p$ and sequence lengths $L$.

spaced-word matches involving the same positions in both sequences), while the second summand comes from background matches. The *relative* weight of the background matches in (4) depends on the match probability $p$ and the sequence length $L$; for $p >> q$ and small $L$, the variance of $N$ is dominated by the 'homologous' term, see Figure 1. Obviously, for large $L$, the background spaced-word matches dominate the 'homologous' ones, since the number of background matches grows quadratically with $L$, while the 'homologous' matches grow only linearly.

Note that, for $L, \ell$ and $w$ fixed, minimizing the variance of $N$ amounts to minimizing

$$\sum_{r \leq r'} \sum_{s \in R(r,r')} p^{n(P_r, P_{r'}, s)} \quad + \quad (L - \ell) \cdot \sum_{r \leq r'} \sum_{s \in R(r,r')} q^{n(P_r, P_{r'}, s)} \qquad (5)$$

Comparison with (2) shows that, in the special case of $p = 1/2$, the first summand of (5) that corresponds to the *homologous* matches is almost identical

with the *overlap complexity* defined by Ilie and Ilie (except for the range $R(r,r)$ in which a pattern $P_r$ is shifted against itself). For sequences of moderate length, the overlap complexity can therefore be seen as an approximation to the variance of the number of spaced-word matches.

In any case, the overlap complexity and the variance of $N$ for a set of pattern $\mathcal{P} = \{P_1, \ldots, P_m\}$ both have the form

$$\sum_{r \leq r'} \alpha_{r,r'}(\mathcal{P}) \tag{6}$$

with

$$\alpha_{r,r'}(\mathcal{P}) = \begin{cases} \displaystyle\sum_{s=1-\ell_{r'}}^{\ell_r - 1} 2^{\sigma_{r,r'}[s]} & (OC) \\ (L-\ell+1) \displaystyle\sum_{s \in R(r,r')} \left( p^{n(P_r,P_{r'},s)} + (L-\ell) \cdot q^{n(P_r,P_{r'},s)} \right) & (Var) \end{cases} \tag{7}$$

Our optimization problem is therefore: for integers $m, \ell_1, \ldots \ell_m, w$, find a set $\mathcal{P}$ of $m$ patterns of lengths $\ell_1, \ldots, \ell_m$ and weight $w$ that minimizes the sum (6).

## Hill-climbing algorithms to find sets of patterns with minimal $Var(N)$ or $OC$

Both *SpEED* and our new algorithm start with a randomly generated pattern set $\mathcal{P}$ and use *hill-climbing* to gradually reduce the $OC$ or $Var(N)$ of $\mathcal{P}$. After a pattern set with low $OC$ is obtained in this way, its *sensitivity* is calculated, if one is looking for a pattern set with maximal sensitivity. This step is omitted in *rasbhari* if $Var(N)$ or $OC$ is to be minimized. The whole procedure is repeated, and the pattern set with the overall highest sensitivity – or lowest variance of $N$ or $OC$, respectively – is returned.

### Original hill-climbing algorithm

To improve the current pattern set $\mathcal{P}$, the hill-climbing algorithm implemented in *SpEED* looks at all triples $(r,i,j)$ where $P_r$ is a pattern in $\mathcal{P}$, and $i$ and $j$ are a *match position* and a *don't-care* position in $P_r$, respectively. For each such triple $(r,i,j)$, the algorithm considers the pattern set that would be obtained from $\mathcal{P}$ by swapping $i$ and $j$ in $P_r$ – *i.e.* by turning $i$ into a don't-care and $j$ into a match position. The $OC$ is calculated for all

pattern sets that can be obtained in this way, and the one with the lowest $OC$ is selected as the next pattern set $\mathcal{P}$. This is repeated iteratively.

There are $O(m \cdot \ell^2)$ triples $(r, i, j)$ to be considered to modify the current pattern set $\mathcal{P}$. For each of these triples, the $OC$ is to be calculated for the pattern set that would be obtained by swapping $i$ and $j$ in $P_r$. To this end, the modified pattern $P_r$ has to be compared to the $m - 1$ remaining patterns in $\mathcal{P}$ which, for each pattern comparison, involves $O(\ell)$ shifts of two patterns against each other. In each shift, the number of common match positions is to be counted, which takes again $O(\ell)$ time. Thus, calculating the $OC$ of the pattern set obtained by swapping two positions $i$ and $j$ in a pattern $P_r$ takes $O(m \cdot \ell^2)$ time, so finding an optimal triple $(r, i, j)$ to determine the next pattern set takes $O(m^2 \cdot \ell^4)$ time. This step is repeated a certain number of times; for the pattern set that is finally obtained by this hill-climbing routine, the *sensitivity* is calculated. This whole procedure is repeated 5,000 times, and finally the set with the best sensitivity is returned.

**Modified hill-climbing algorithm**

In our modified hill-climbing algorithm, we also swap a match position $i$ with a don't-care position $j$ in some pattern $P_r$ in each step of the algorithm, and we evaluate the pattern set that would be obtained by this operation. However, instead of looking at *all* possible triples $(r, i, j)$, we look at those patterns first that contribute most to the $OC$ or $Var(N)$, respectively, of the current pattern set $\mathcal{P}$. The contribution

$$C_r \quad = \quad \sum_{r'} \alpha_{r,r'} \tag{8}$$

of a pattern $P_r \in \mathcal{P}$ to the $OC$ or $Var(N)$ of $\mathcal{P}$ can be calculated as a by-product, whenever $OC$ or $Var(N)$ is calculated, with $\alpha$ as in (7). We then sort the patterns in $P_r \in \mathcal{P}$ according to the values $C_r$, and we process them in descending order of $C_r$, *i.e.* we look at those patterns first that contribute *most* to the $OC$ or $Var(N)$ of $\mathcal{P}$.

For the current pattern in the list, we randomly select a match position $i$ and a don't-care position $j$. If swapping $i$ and $j$ does *not* improve the current pattern set, we move on to the next pattern in the list and proceed in the same way. This is repeated until we find a pattern where swapping the selected pair of random positions does improve $\mathcal{P}$. In this case, the modified pattern is accepted, all values $C_r$ are updated, the patterns in $\mathcal{P}$ are sorted accordingly, and we start again with the pattern $P_r$ with maximum $C_r$. If we reach the last pattern in the list without obtaining any improvement, we

start again with the first pattern, *i.e.* the pattern with the largest $C_r$, select new random positions $i$ and $j$ etc. Processing one pattern $P_r$ in this way takes $O(m \cdot \ell^2)$ time, since we look only at one single pair $(i, j)$ and calculate the $OC$ or $Var(N)$ of the pattern set that would be obtained by swapping $i$ and $j$ in $P_r$.

The hill climbing is continued until a user-defined number of triples $(r, i, j)$ have been processed, then the current pattern set is returned; by default, 25,000 triples are processed. If we want to obtain a pattern set with maximal *sensitivity*, the described hill-climbing procedure is repeated 100 times, and for the pattern set with the lowest $OC$ among the 100 obtained pattern sets, the *sensitivity* is calculated. To calculate the sensitivity, *rasbhari* uses program code from *SpEED*. Again, this whole process is repeated 5,000 times, so for a total of 5,000 pattern sets the sensitivity is calculated during one program run. This is similar to *SpEED*, but in *SpEED* the time-consuming sensitivity calculation is carried out after *one* round of hill climbing. By contrast, we run our faster hill-climbing routine 100 times before we calculate the sensitivity for the *best* pattern set from these 100 runs. The final output of our program is the pattern set with the highest sensitivity from the 5,000 iterations.

The number $m$ of patterns and their weight $w$ are to be specified by the user. If $Var(N)$ is to be minimized for alignment-free sequence comparison, all patterns have the same length $\ell$ which is also to be specified by the user. If the *sensitivity* is to be maximized for database searching and read alignment, better results are achieved if the patterns in $\mathcal{P}$ have different lengths. In this case, the maximum and minimum pattern lengths need to be specified. The program then selects lengths $\ell_1, \ldots, \ell_m$ that are evenly distributed between these extreme values.

# Results

## Sensitivity in database searching

To evaluate *rasbhari*, we first applied it to generate pattern sets, maximizing the *sensitivity* for database searching and read mapping. For the number $m$ and weight $w$ of the patterns and for the length $H$ and match probability $p$ of the homology regions, we used the parameter settings from *SHRiMP2* [12], *PatternHunter II* [27] and *BFAST* [18]. We and compared it to the sensitivity of pattern sets obtained with *Iedera* [23], *SpEED* [21], *AcoSeeD* [14], *FastHC* and *MuteHC* [13] as published by the authors of these programs; the results of this comparison are shown in Table 1. Here, the sensitivity

Table 1: Sensitivity of pattern sets in hit-and-extend database searching, calculated from different programs. Parameter settings for the number $m$ and weight $w$ of patterns, the length $H$ of the gap-free homology region between query and database sequences and the match probability $p$ in the homology regions, are taken from three popular programs *SHRiMP2, PatternHunter II* and *BFAST*. Results of existing programs are taken from their respective publications.

| $w$ | $p$ | *Iedera* | *SpEED* | *AcoSeeD* | *FastHC* | *MuteHC* | *rasbhari* |
|---|---|---|---|---|---|---|---|
| ***SHRiMP2:* 4 patterns ($H = 50$)** | | | | | | | |
| 10 | 0.75 | 90.6820 | 90.9098 | 90.9513 | 90.7312 | **92.6812** | 90.9614 |
| | 0.80 | 97.7586 | 97.8337 | 97.8521 | 97.7625 | **98.3836** | 97.8554 |
| | 0.85 | 99.7437 | 99.7569 | 99.7614 | 99.7431 | **99.8356** | 99.7618 |
| 11 | 0.75 | 83.2413 | 83.3793 | **83.4728** | 83.3068 | 83.4127 | 83.4679 |
| | 0.80 | 94.9350 | 94.9861 | 95.037 | 94.9453 | 95.0194 | **95.0386** |
| | 0.85 | 99.2189 | 99.2431 | 99.2478 | 99.2250 | 99.2486 | **99.2506** |
| 12 | 0.80 | 90.3934 | 90.5750 | 90.6328 | 90.4735 | 90.5820 | **90.6648** |
| | 0.85 | 98.0781 | 98.1589 | 98.1766 | 98.1199 | 98.1670 | **98.1824** |
| | 0.90 | 99.8773 | 99.8821 | 99.8853 | 99.8771 | 99.8836 | **99.8864** |
| 16 | 0.85 | 84.5795 | 84.8212 | **84.9829** | 84.6558 | 84.8764 | 84.969 |
| | 0.90 | 97.2806 | 97.4321 | 97.4712 | 97.3556 | 97.4460 | **97.5035** |
| | 0.95 | 99.9331 | 99.9388 | 99.9419 | 99.9347 | 99.9424 | **99.9441** |
| 18 | 0.85 | 72.1695 | 73.1664 | **73.27** | 72.9558 | | 73.2209 |
| | 0.90 | 93.0442 | 93.7120 | 93.7778 | 93.6030 | | **93.78** |
| | 0.95 | 99.6690 | 99.7500 | **99.7599** | 99.7399 | | 99.7557 |
| ***PatternHunterII:* 16 patterns ($H = 64$)** | | | | | | | |
| 11 | 0.70 | 92.0708 | 93.2526 | | 93.0585 | | **93.4653** |
| | 0.75 | 98.3391 | 98.6882 | | 98.6352 | | **98.7573** |
| | 0.80 | 99.8366 | 99.8820 | | 99.8750 | | **99.8907** |
| ***BFAST:* 10 patterns ($H = 50$)** | | | | | | | |
| 22 | 0.85 | 60.1535 | 60.8127 | | 60.0943 | | **60.9919** |
| | 0.90 | 87.9894 | 88.5969 | | 88.0426 | | **88.8005** |
| | 0.95 | 99.2196 | 99.3659 | | 99.2923 | | **99.4099** |

values of *AcoSeeD* are *average* values over 10 program runs reported in [14]. If the sensitivity of a pattern set is to be optimized, the run time of *rasbhari* is comparable to *SpEED*, since the most time-consuming step in both programs is to calculate the sensitivity of a current pattern set $\mathcal{P}$ which is done

5,000 times per program run in each of the two programs.

### Alignment-free phylogeny reconstruction

Next, we wanted to know how alignment-free phylogeny reconstruction can be improved with *rasbhari*. To this end, we estimated pairwise distances between simulated DNA sequences using the *Spaced Words* approach described in [34], and we measured the accuracy of these distance estimates for different underlying pattern sets. Here, we used *rasbhari* to minimize the *variance* of the number $N$ of spaced-word matches between two sequences. Since there is no other method to minimize the variance of $N$, we compared the pattern sets from *rasbhari* with the randomly generated pattern sets that we initially used. The phylogenetic distances obtained from both program runs were compared to the 'real' distances, *i.e.* the average number of substitutions per position in the sequences.

Table 2: Accuracy of pattern sets in alignment-free phylogeny reconstruction. Evolutionary distances between simulated DNA sequences were estimated based on the number $N$ of spaced-word matches between them, using the method published in [34]. Four sets of sequence pairs were generated, such that for each set the substitution rates of the sequence pairs are within a certain range, as specified in the table. Pattern sets were generated with *rasbhari*, minimizing the variance of $N$, and compared to the randomly generated pattern sets that we previously used. The table shows the mean quadratic difference between the estimated distances and the 'real' distances, *i.e.* the number of substitutions per position in the model used to generate the sequence pairs.

| Substitutions per position | 0 - 0.25 | 0.25 - 0.5 | 0.5 - 0.75 | 0.75 - 1.0 |
|---|---|---|---|---|
| Random pattern sets | 1.86E-06 | 5.10E-05 | 1.92E-03 | 3.11E-02 |
| *rasbhari* | 1.69E-06 | 4.51E-05 | 1.60E-03 | 3.01E-02 |

As test data, we generated four data sets with 10,000 pairs of DNA sequences of length 100 *kb* each. In each of these data sets, we used a specific range of substitution rates as shown in Table 2; the substitution rate for a simulated sequence pair was then randomly picked within the respective range. For each program run, we used a set of $m = 3$ patterns of length 20 with 16 *match* and 4 *don't-care* positions. Table 2 shows the *mean quadratic difference* between the estimated distances and the 'real' distance between the sequences. In all four categories, the pattern sets generated

with *rasbhari* were superior to the randomly generated pattern sets.

### Read classification

As a third test case, we used different pattern sets for *CLARK-S* [38], a recently developed tool for short read classification. We evaluated the accuracy of *CLARK-S* with three underlying pattern sets, namely **(A)** with the patterns used by default in the program, **(B)** with patterns from *rasbhari* minimizing *overlap complexity* and **(C)** with patterns from *rasbhari* maximizing *sensitivity*. *CLARK-S* uses sets of $m = 3$ patterns of length $\ell = 31$ and with a weight of $w = 22$. Since *SpEED* is too slow to generate pattern sets for these parameters, the authors generated pattern sets for *CLARK-S* by exhaustively searching over all *single* patterns with $\ell = 31$ and $w = 22$. If the first and the last position in the patterns are required to be *match positions*, this approach has to evaluate $\binom{29}{20} \approx 10^7$ possible patterns. The sensitivity of each of these patterns was calculated, and the three patterns with the highest sensitivity were selected. Note however, that maximizing the sensitivity of *single* patterns is only an approximation to finding a *set* of patterns with maximal *total* sensitivity.

Figure 2 shows the default pattern set from *CLARK-S* and the two pattern sets generated by *rasbhari* as described. The exhaustive procedure used by *CLARK-S* took 2 hours to generate the pattern set. *rasbhari*, by contrast, calculated pattern sets with the same parameters within 7.54 seconds with the *slow* version where the *sensitivity* is calculated, and within 0.068 seconds with the *fast* version where the *overlap complexity* is maximized without considering the sensitivity explicitly. The slow version of *rasbhari* is thus around 480 times faster than the exhaustive procedure in *CLARK-S*, while the fast version is around $52,000$ times faster. The *theoretical* sensitivity of the three pattern sets is 0.999771 for the default patterns from *CLARK-S*, 0.999811 for the *rasbhari* patterns with minimized overlap complexity and 0.999822 for the *rasbhari* patterns with maximized sensitivity.

To evaluate the classification accuracy of *CLARK-S* with these three pattern sets experimentally, we used five data sets from the literature, namely two sets, *HC1* and *HC2*, from the *MetaPhlAn* project [42] and three sets, *simHC, simMC and simLC*, from the *FAMeS* databases [31]. For each of these data sets, we calculated *precision* and *sensitivity* of the classification at the *species level* as defined in [39]. That is, for a classification task where objects are to be assigned to classes, *precision* is defined as the fraction of correct assignments among the total number of assignments, while *sensitivity* is the ratio between the number of correct assignments and the number

**(A)**

```
111110110100111010011101101111
111011101110010111001011011111
111110010111011010110011011111
```

**(B)**

```
111101011101001100111011110111
111011101110111101001011001111
111110100101110011111010110111
```

**(C)**

```
111011110011010111110101011011
111010101110110011010011111111
111111010110101110011011001111
```

Figure 2: Pattern sets used for short read classification: **(A)** as used by default in *CLARK-S*, **(B)** generated with *rasbhari* minimizing *overlap complexity* and **(C)** generated with *rasbhari* maximizing *sensitivity*.

of objects to be classified. The two values are not the same since not every object is necessarily assigned to one of the classes; *precision* is always larger than or equal to *sensitivity* since the denominator in the definition of precision is smaller or equal to the denominator in the definition of sensitivity. Since this definition of *sensitivity* refers to the ability of a program to correctly classify objects, it is not to be confused with the sensitivity in database searching as discussed above. Table 3 summarizes precision and sensitivity of *CLARK-S* with its default pattern set and with a pattern set generated by *rasbhari*.

Table 3: Read classification with *CLARK-S* [38] with the default pattern set of the program and with the pattern set from *rasbhari* for the same parameter values, namely $n = 3$ patterns of length $\ell = 31$ and weight $w = 21$. *Precision* and *sensitivity* of the classification are reported at the *species level* for five data sets from the literature.

| | | Default pattern set | | rasbhari | |
|---|---|---|---|---|---|
| Dataset | #reads | Precision | Sensitivity | Precision | Sensitivity |
| HC1 | 999,998 | 97.69 | 90.36 | 97.69 | **90.44** |
| HC2 | 999,991 | 96.45 | 88.11 | 96.45 | **88.18** |
| simHC | 116,771 | 97.20 | 90.53 | 97.20 | **90.54** |
| simMC | 97,495 | **98.75** | 95.09 | 98.73 | 95.09 |
| simLC | 114,457 | **98.29** | **94.26** | 98.28 | 94.25 |

### Improvement of *OC* by hill climbing

Figure 3 shows how the overlap complexity *(OC)* of pattern sets produced by *rasbhari* depends on the number of iteration steps carried out in the hill-climbing algorithm. For a set of $m = 10$ patterns of length $\ell = 14$ and weight $w = 8$, a single run of the hill-climbing procedure converges after around 3,000 steps; for $m = 20, \ell = 44, w = 14$, it converges after around 80,000 steps. The *OC* is further improved if the hill-climbing procedure is run multiple times and the best result of these runs is used.

## Discussion

We developed a program called *rasbhari* to calculate sets of binary *patterns* – or *spaced seeds*, as they are often called – for read mapping, database searching and alignment-free sequence comparison. For sequence-homology searching, *rasbhari* optimizes the *sensitivity* of pattern sets, *i.e.* the probability of obtaining at least one hit between a query and a database sequence that share a gap-free homology of a given length and with a given match probability between nucleotides. Since the sensitivity of a pattern set is expensive to calculate, our algorithm optimizes the *overlap complexity* of the produced pattern sets which is closely related to its sensitivity. We use a *hill-climbing* algorithm, similar to the one used in *SpEED*, to minimize the overlap complexity. Unlike *SpEED*, however, our algorithm does not calculate the overlap complexity of *all* neighbours of a current pattern set, but modifies those patterns first that contribute most to the overlap complexity

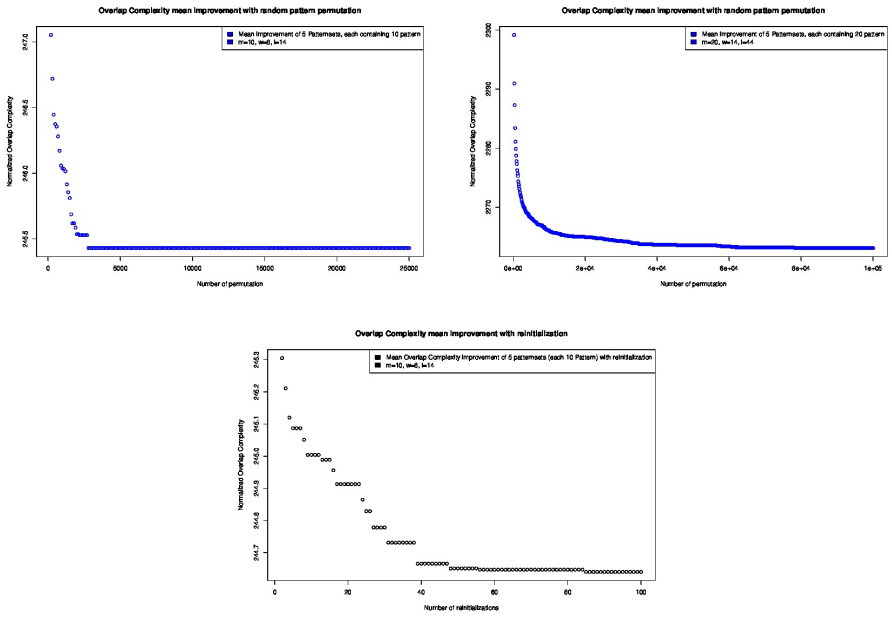
Overlap Complexity mean improvement with random pattern permutation
Mean improvement of 5 Patternsets, each containing 10 pattern
m=10, w=8, l=14
Normalized Overlap Complexity
Number of permutation
Overlap Complexity mean improvement with random pattern permutation
Mean improvement of 5 Patternsets, each containing 20 pattern
m=20, w=14, l=44
Normalized Overlap Complexity
Number of permutation
Overlap Complexity mean improvement with reinitialization
Mean Overlap Complexity improvement of 5 patternsets (each 10 Pattern) with reinitialization
m=10, w=8, l=14
Normalized Overlap Complexity
Number of reinitializations

Figure 3: Normalized overlap complexity *(OC)* of pattern sets depending on the number of iteration steps in our algorithm. The first two plots show how the $OC$ is reduced in a single round of the hill-climbing algorithm for different parameters. For a set of $m = 10$ patterns of length $\ell = 14$ and weight $w = 8$, the algorithm converges after around 3,000 iteration steps of hill-climbing (upper plot); for a set of $m = 20$ patterns of length $\ell = 44$ and weigth $w = 14$, it converges after around 80,000 steps (middle plot). The lower plot shows how the $OC$ is improved if the hill-climbing algorithm is run multiple times and the best result of all runs is returned.

of the current pattern set. If maximizing the *sensitivity* in database searching, we calculate the sensitivity of the current pattern set after a certain number of iterations and, finally, the pattern set with the overall highest sensitivity is returned.

As a fast alternative, *rasbhari* can minimize the overlap complexity alone, without calculating the sensitivity of pattern sets. This option is useful in situations where large pattern sets are needed for which it would take too long to calculate the sensitivity. As a third option, *rasbhari* can minimize the variance of the number $N$ of spaced-word matches in alignment-free

sequence comparison which is used by various methods to estimate phylogenetic distances between sequences. We could show that, mathematically, the variance of $N$ has a similar form as the overlap complexity of a pattern set, so the same optimization algorithm can be used to minimize both of them.

In both homology searching and read classification, pattern sets generated by *rasbhari* are more *sensitive* than the default pattern sets, so more homologies can be detected and more reads can be correctly classified. At first glance, the increase in sensitivity that we obtained seems moderate; as shown in Table 1, the improvement is usually in the first or second digit after the decimal mark. In database searching and read mapping, however, even small improvements in sensitivity can lead to a large number of additional hits. Moreover, as these additional hits will be mostly in the 'twilight zone' of low sequence similarity, they may be of particular interest to the user.

In the context of read alignment, Ilie *et al.* pointed out that, with a 100-fold coverage of the human genome, a 1 percent improvement in pattern sensitivity would mean that 3 billion more nucleotides could be mapped [21], so the improvement that we achieved with *rasbhari* would still lead to tens or hundreds of millions of additionally mapped nucleotides. In database searching, the situation is similar. If we consider, for example, homology regions of length $H = 64$ with a match probability of $p = 0.8$ at the nucleotide level, then with $w = 11$, the sensitivity of *rasbhari* is improved by less than 0.01 percentage points compared to *SpEED*, see Table 1. Note, however, that these sensitivity values are already close to 100%, so the fraction of homologies that are *not* detected can be considerably reduced with the sligh improvement in sensitivity obtained with *rasbhari*. In our example, the number of homologies that are *missed* is reduced by $> 7\%$ if *rasbhari* is used instead of *SpEED*. With the same parameters, but with $p = 0.7$, the sensitivity of both programs is around 93%. Here, the number of missed homologies is still reduced by 3% with *rasbhari*, compared to *SpEED*.

For alignment-free sequence comparison, pattern sets produced by *rasbhari* lead to more accurate phylogenetic distances than the random pattern sets that we previously used. While this result may not be surprising, *rasbhari* is, to our knowledge, the first program that has been designed for this purpose and that can minimize the variance of the number of spaced-word matches. We therefore integrated *rasbhari* into our web server for alignment-free sequence comparison [19].

In read classification, the sensitivity of *CLARK-S* could be increased by 0.08 and $0,07$ percentage points, respectively, for the largest data sets that we used, *HC1* and *HC2*. Each of these data sets contains around one million

reads, so the improvement in sensitivity that we achieved with *rasbhari* means that 800 more reads from *HC1* and 700 more from *HC2* could be correctly classified by *CLARK-S*. This is remarkable, since the classification accuracy of *CLARK-S* is already very high, so it is hard to further improve the program. An interesting question in the context of *CLARK-S* is how the length and weight of the patterns influence its accuracy. So far, it was difficult to investigate this question systematically, since the exhaustive method that the program uses by default, is too time consuming. With the massive improvement in runtime that we achieved with *rasbhari*, it is now possible to systematically investigate how the accuracy of *CLARK-S* depends on the parameters of the underlying pattern sets.

In the hill-climbing procedure, our default of 25,000 iteration steps was sufficient to obtain stable results for the parameter settings that we used in our benchmark studies; we were unable to further improve these results by increasing the number of iterations. For different values of $m, w, \ell, p$ and $H$, however, it may be advisable to adapt the number of iteration steps. If the number of patterns or their length and weight are increased, Figure 3 shows that a larger number of iteration steps can improve the results. The number of iterations within one round of hill climbing and the number of times the hill-climbing is carried out can be passed to *rasbhari* through the command line; the users can therefore adapt these parameter values for their particular applications if they do not want to use the default values of the program.

## Author Contribution:

LH conceived and implemented the algorithm, CL evaluated *rasbhari* for phylogeny reconstruction, RO and SL evaluated the accuracy of read classification with *CLARK-S*, BM guided the study and wrote the manuscript. All authors read and approved the manuscript.

## Acknowledgements

We would like to thank Laurent Noé for helpful discussions and for pointing out the similarity between the overlap complexity and the variance of the number of spaced-word matches. Lucian Ilie made useful comments on a previous version of this paper.

## References


[1] E. S. Allman, J. A. Rhodes, and S. Sullivant. Statistically-consistent $k$-mer methods for phylogenetic tree reconstruction. *arXiv:1511.01956 [q-bio.PE]*.

[2] S. F. Altschul, W. Gish, W. Miller, E. M. Myers, and D. J. Lipman. Basic local alignment search tool. *Journal of Molecular Biology*, 215:403–410, 1990.

[3] G. Bernard, C. X. Chan, and M. A. Ragan. Alignment-free microbial phylogenomics under scenarios of sequence divergence, genome rearrangement and lateral genetic transfer. *Scientific Reports*, 6:28970, 2016.

[4] B. Brejova, D. G. Brown, and T. Vinar. Optimal spaced seeds for homologous coding regions. *Journal of Bioinformatics and Computational Biology*, 1:595–610, 2004. Early version appeared in CPM 2003.

[5] B. Brejova, D. G. Brown, and T. Vinar. Vector seeds: an extension to spaced seeds. *Journal of Computer and System Sciences*, 70:364–380, 2005.

[6] B. Buchfink, C. Xie, and D. H. Huson. Fast and sensitive protein alignment using DIAMOND. *Nature Methods*, 12:59–60, 2015.

[7] K. Břinda, M. Sykulski, and G. Kucherov. Spaced seeds improve $k$-mer-based metagenomic classification. *Bioinformatics*, 31:3584–3592, 2015.

[8] A. Califano and I. Rigoutsos. FLASH: a fast look-up algorithm for string homology. In *Computer Vision and Pattern Recognition, 1993. Proceedings CVPR '93., 1993 IEEE Computer Society Conference on*, pages 353–359, 1993.

[9] B. Chor, D. Horn, Y. Levy, N. Goldman, and T. Massingham. Genomic DNA $k$-mer spectra: models and modalities. *Genome Biology*, 10:R108, 2009.

[10] A. E. Darling, B. Mau, and N. T. Perna. progressiveMauve: Multiple Genome Alignment with Gene Gain, Loss and Rearrangement. *PLOS ONE*, 5:e11147+, 2010.

[11] A. E. Darling, T. J. Treangen, L. Zhang, C. Kuiken, X. Messeguer, and N. T. Perna. *Algorithms in Bioinformatics: 6th International Workshop, WABI 2006, Zurich, Switzerland, September 11-13, 2006. Proceedings*, chapter Procrastination Leads to Efficient Filtration for Local Multiple Alignment, pages 126–137. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[12] M. David, M. Dzamba, D. Lister, L. Ilie, and M. Brudno. SHRiMP2: Sensitive yet practical short read mapping. *Bioinformatics*, 27:1011–1012, 2011.

[13] P.-T. Do and C.-G. Tran-Thi. An improvement of the overlap complexity in the spaced seed searching problem between genomic dnas. In *Proceedings of the 2nd National Foundation for Science and Technology Development Conference on Information and Computer Science (NICS)*, pages 271–276, 2015.

[14] D. D. Duc, H. Q. Dinh, T. H. Dang, K. Laukens, and H. H. Xuan. AcoSeeD: An ant colony optimization for finding optimal spaced seeds in biological sequence search. In *ANTS'12*, pages 204–211, 2012.

[15] M. C. Frith and L. Noé. Improved search heuristics find 20 000 new alignments between human and mouse genomes. *Nucleic Acids Research*, 42:e59, 2014.

[16] B. Haubold, N. Pierstorff, F. Möller, and T. Wiehe. Genome comparison without alignment using shortest unique substrings. *BMC Bioinformatics*, 6:123, 2005.

[17] H. Hauswedell, J. Singer, and K. Reinert. Lambda: the local aligner for massive biological data. *Bioinformatics*, 30:i349–i355, 2014.

[18] N. Homer, B. Merriman, and S. F. Nelson. Bfast: an alignment tool for large scale genome resequencing. *PLOS ONE*, 4:e7767+, 2009.

[19] S. Horwege, S. Lindner, M. Boden, K. Hatje, M. Kollmar, C.-A. Leimeister, and B. Morgenstern. *Spaced words* and *kmacs*: fast alignment-free sequence comparison based on inexact word matches. *Nucleic Acids Research*, 42:W7–W11, 2014.

[20] L. Ilie and S. Ilie. Multiple spaced seeds for homology search. *Bioinformatics*, 23:2969–2977, 2007.

[21] L. Ilie, S. Ilie, and A. M. Bigvand. SpEED: fast computation of sensitive spaced seeds. *Bioinformatics*, 27:2433–2434, 2011.

[22] S.-R. Jun, G. E. Sims, G. A. Wu, and S.-H. Kim. Whole-proteome phylogeny of prokaryotes by feature frequency profiles: An alignment-free method with optimal feature resolution. *Proceedings of the National Academy of Sciences*, 107:133–138, 2010.

[23] G. Kucherov, L. Noé, and M. Roytberg. A unifying framework for seed sensitivity and its application to subset seeds. *Journal of Bioinformatics and Computational Biology*, 4:553–569, 2006.

[24] C.-A. Leimeister, M. Boden, S. Horwege, S. Lindner, and B. Morgenstern. Fast alignment-free sequence comparison using spaced-word frequencies. *Bioinformatics*, 30:1991–1999, 2014.

[25] C.-A. Leimeister and B. Morgenstern. *kmacs*: the $k$-mismatch average common substring approach to alignment-free sequence comparison. *Bioinformatics*, 30:2000–2008, 2014.

[26] C. S. Leslie, E. Eskin, A. Cohen, J. Weston, and W. S. Noble. Mismatch string kernels for discriminative protein classification. *Bioinformatics*, 20:467–476, 2004.

[27] M. Li, B. Ma, D. Kisman, and J. Tromp. PatternHunter II: highly sensitive and fast homology search. *Journal of Bioinformatics and Computational Biology*, 02:417–439, 2004.

[28] M. Li, B. Ma, and L. Zhang. Superiority and complexity of the spaced seeds. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA '06, pages 444–453, Philadelphia, PA, USA, 2006. Society for Industrial and Applied Mathematics.

[29] T. Lingner and P. Meinicke. Remote homology detection based on oligomer distances. *Bioinformatics*, 22:2224–2231, 2006.

[30] B. Ma, J. Tromp, and M. Li. PatternHunter: faster and more sensitive homology search. *Bioinformatics*, 18:440–445, 2002.

[31] K. Mavromatis, N. Ivanova, K. Barry, H. Shapiro, E. Goltsman, A. C. McHardy, I. Rigoutsos, A. Salamov, F. Korzeniewski, M. Land, A. Lapidus, I. Grigoriev, P. Richardson, P. Hugenholtz, and N. C. Kyrpides. Use of simulated data sets to evaluate the fidelity of metagenomic processing methods. *Nature Methods*, 4:495–500, 2007.

[32] P. Meinicke. UProC: tools for ultra-fast protein domain classification. *Bioinformatics*, 31:1382–1388, 2015.

[33] P. Meinicke, M. Tech, B. Morgenstern, and R. Merkl. Oligo kernels for datamining on biological sequences: a case study on prokaryotic translation initiation sites. *BMC Bioinformatics*, 5:169, 2004.

[34] B. Morgenstern, B. Zhu, S. Horwege, and C.-A. Leimeister. Estimating evolutionary distances between genomic sequences from spaced-word matches. *Algorithms for Molecular Biology*, 10:5, 2015.

[35] L. Noé, M. Gîrdea, and G. Kucherov. Designing efficient spaced seeds for SOLiD read mapping. *Advances in Bioinformatics*, 2010:1–12, 2010.

[36] L. Noé and D. E. K. Martin. A coverage criterion for spaced seeds and its applications to SVM string-kernels and $k$-mer distances. *Journal of Computational Biology*, 12:947–963, 2014.

[37] T. Onodera and T. Shibuya. The gapped spectrum kernel for support vector machines. In P. Perner, editor, *Machine Learning and Data Mining in Pattern Recognition*, volume 7988 of *Lecture Notes in Computer Science*, pages 1–15. Springer Berlin Heidelberg, Berlin,Heidelberg, 2013.

[38] R. Ounit and S. Lonardi. *Algorithms in Bioinformatics: 15th International Workshop, WABI 2015, Atlanta, GA, USA, September 10-12, 2015, Proceedings*, chapter Higher Classification Accuracy of Short Metagenomic Reads by Discriminative Spaced $k$-mers, pages 286–295. Springer Berlin Heidelberg, Berlin, Heidelberg, 2015.

[39] R. Ounit, S. Wanamaker, T. J. Close, and S. Lonardi. Clark: fast and accurate classification of metagenomic and genomic sequences using discriminative $k$-mers. *BMC Genomics*, 16:236, 2015.

[40] S. M. Rumble, P. Lacroute, A. V. Dalca, M. Fiume, A. Sidow, and M. Brudno. SHRiMP: Accurate mapping of short color-space reads. *PLOS Computational Biology*, 5:e1000386+, 2009.

[41] S. Schbath, V. Martin, M. Zytnicki, J. Fayolle, V. Loux, and J.-F. Gibrat. Mapping reads on a genomic sequence: An algorithmic overview and a practical comparative analysis. *Journal of Computational Biology*, 19:796–813, 2012.

[42] N. Segata, L. Waldron, A. Ballarini, V. Narasimhan, O. Jousson, and C. Huttenhower. Metagenomic microbial community profiling using unique clade-specific marker genes. *Nature Methods*, 9:811–814, 2012.

[43] G. E. Sims, S.-R. Jun, G. A. Wu, and S.-H. Kim. Alignment-free genome comparison with feature frequency profiles (FFP) and optimal resolutions. *Proceedings of the National Academy of Sciences*, 106:2677–2682, 2009.

[44] S. V. Thankachan, A. Apostolico, and S. Aluru. A provably efficient algorithm for the $k$-mismatch average common substring problem. *Journal of Computational Biology*, 2016.

[45] S. V. Thankachan, S. P. Chockalingam, Y. Liu, A. Krishnan, and S. Aluru. A greedy alignment-free distance estimator for phylogenetic inference (extended abstract). In *Computational Advances in Bio and Medical Sciences (ICCABS), 2015 IEEE 5th International Conference on*, pages 1–1, Oct 2015.

[46] S. V. Thankachan, S. P.Chockalingam, Y. Liu, A. Apostolico, and S. Aluru. ALFRED: a practical method for alignment-free distance computation. *Journal of Computational Biology,* in press.

[47] I. Ulitsky, D. Burstein, T. Tuller, and B. Chor. The average common substring approach to phylogenomic reconstruction. *Journal of Computational Biology*, 13:336–350, 2006.

[48] S. Vinga. Editorial: Alignment-free methods in computational biology. *Briefings in Bioinformatics*, 15:341–342, 2014.

[49] S. Vinga, A. M. Carvalho, A. P. Francisco, L. M. S. Russo, and J. S. Almeida. Pattern matching through chaos game representation: bridging numerical and discrete data structures for biological sequence analysis. *Algorithms for Molecular Biology*, 7:10, 2012.